High-Throughput Wireless LANs

Intel research into wireless LAN air interfaces identifies key technologies for dramatic improvements in throughput

Research & Development at Intel

Boyd Bangerter, Eric Jacobsen, Minnie Ho, Adrian Stephens, Alexander Maltsev, Alexey Rubtsov, Ali Sadri, Larry Swanson

Intel's long-standing commitment to wireless LAN (WLAN) technology led it to be one of the first microprocessor vendors to embed Wi-Fi technology in a platform chipset—the Intel® Centrino™ mobile technology released in early 2003. Intel Research is now working to develop and test candidate technologies for higher throughput WLANs. This paper highlights some of these technologies for unlicensed band wireless air interfaces.

Introduction

WLAN technology is one of the fastest-growing segments in the computing and communications equipment market. This growth is the result of coordinated standards efforts as well as innovation by leading-edge companies that are improving antenna systems, radio frequency (RF) components, modulation schemes, medium access control (MAC) mechanisms, and security mechanisms. Today, the air interface of standards-based WLAN equipment supports data rates up to 54 Mbps. Designers of WLAN equipment are attempting to specify a future air interface that will achieve data rates in excess of 250 Mbps.

The Challenges of High-Speed Design in Unlicensed Band Wireless Channels

Wireless channels present challenging engineering problems not encountered in wireline systems. Most WLAN systems use omnidirectional antennas that provide good coverage but do not concentrate the transmitted power to the intended user. This design increases the amount of signal energy that is scattered and reflected from objects in the environment; and as a result, components of the signal arriving at the receiver are spread out over a longer time period than is desirable. In addition, because the frequency spectrum used by WLAN devices is unlicensed, other devices may attempt to use the same channel resources and thus create interference. The challenge is then to provide a high-performance, reliable data link that maintains high throughput despite restricted receiver power levels, severe channel fading due to multipath reflections, and interference from other devices.

High-Throughput Wireless LAN Candidate Technologies

Fortunately, there are technologies that can take advantage of multipath in wireless channels. One of these, called multiple-input multiple-output (MIMO), uses multiple antennas for both transmitting and receiving. (Figure 1) Another set of technologies uses wideband adaptive orthogonal frequency division multiplexing (OFDM). The following sections describe both sets of technologies.

MIMO Systems for High Throughput

The primary advantage of MIMO systems is that the throughput capacity of the system (with equal numbers of transmit and receive antenna elements) increases almost directly in proportion to the number of transceiver antenna elements without increasing the total transmitted power or frequency bandwidth. This capacity is achieved by dividing the channel into multiple spatial channels through which independent data streams can be transmitted. This technique is known as spatial multiplexing.

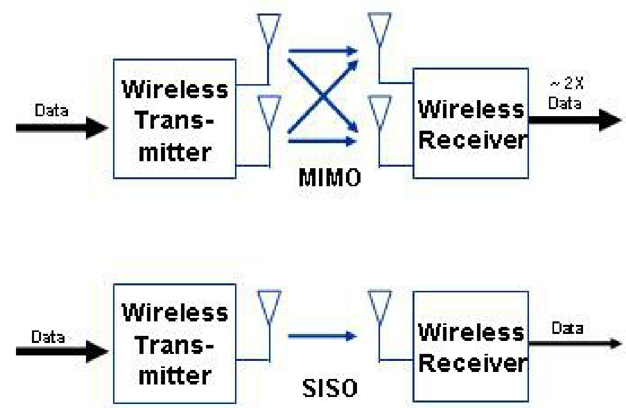

Figure 1. The differences between MIMO and systems that use single-input, single output (SISO).

Additionally, a MIMO system can be used to increase range by increasing diversity gain. One way to do this is to use a transmitter to simply encode the bits over space and over frequency and transmit them over multiple spatial channels. The receiver then separates the symbols from the multiple spatial channels and decodes the bits. Another approach requires the transmitter to have some knowledge of the channel, which is obtained by feedback from the receiver. The transmitter then uses this knowledge of the channel to increase diversity gain. Simulations have shown that by doubling the number of antennas and using channel knowledge to adjust diversity gain, bandwidth throughput can increase by 100% and range can increase up to 40% over what is conventional today.

Wideband Adaptive OFDM

OFDM is the modulation currently used by 802.11a and 802.11g systems. Much research has gone into achieving the theoretical maximum channel capacity using adaptive modulation, subcarrier power allocation, and coding techniques for OFDM systems. The underlying idea behind these methods involves dividing a wideband channel into a larger number of sub-channels each of which can contain

several subcarriers. The concept is then to put more power and higher-order modulations onto the subcarriers with larger signal-to-noise ratios (SNR), while lower SNR subcarriers receive less power and lower-order modulation to a certain threshold, after which the subcarrier is simply turned off. In other words, the idea is to exploit "good" and "bad" subcarriers in such a way as to maximize the data-carrying ability of the channel. This concept is illustrated in Figure 2.

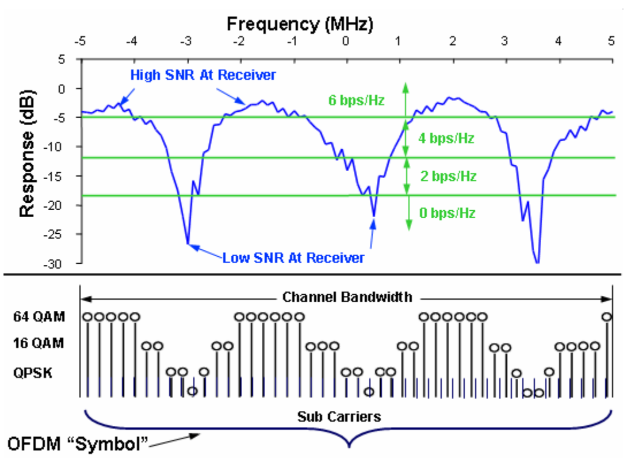

Figure 2. Adaptive bit and power loading based on channel SNR.

As the channel varies over time, further adaptations can be made in each subchannel to continually optimize the data-carrying capacity of the channel. Intel has been researching different types of adaptive bit and power-loading schemes that guarantee a target error rate, but differ in throughput performance, complexity, and the amount of feedback service information. Three schemes are described next.

Adaptive Bit and Power-Loading (ABPL)

With ABPL, after an initial measurement of the channel, the number of subcarriers in each subset with a uniform modulation plus encoding type is determined. SNR pre-equalization is then performed within each subset in order to guarantee a target fixed bit error rate (BER) for each subcarrier. During this procedure, surplus power that is removed from higher-order modulation subsets is added to lower-order modulation subsets, and it is also used to turn on subcarriers that were initially turned off. ABPL is the most computationally complicated scheme, but it has the best throughput performance (approximately 80% better than conventional OFDM). However, it requires the most feedback service information, and it is very sensitive to channel state estimation errors.

Adaptive Bit-Loading (ABL) Scheme

This scheme uses the same transmit power for all active subcarriers and uses a "no transmission" mode for very bad subcarriers. The ABL scheme has some throughput degradation compared to ABPL, but it needs less feedback service information. Simulations have shown that ABL is more robust to impairments such as inaccurate channel state estimation and to interference and time variation of the wireless channel.

To decrease the amount of channel feedback required, Intel has investigated modifications of these schemes that allocate subcarriers of the same parameters in groups of 2, 4, or 6 subcarriers. In each group, subcarriers have the same power and modulation/coding type. Simulations show that a 2-subcarrier grouping has negligible throughput performance degradation and it halves the amount of feedback information required, whereas groupings of 4 and 6 lead to performance degradation.

Optimized Modulation and Coding Selection (MCS) Scheme

Another fast adaptation scheme gives all subcarriers the same modulation and power, based on channel state feedback. The optimal selection of modulation and coding type is made on the basis of the "momentary link performance," that is., short-term SNR or packet error rate (PER) estimates. Intel has investigated algorithms for an MCS scheme. While MCS shows throughput degradation compared with ABPL or ABL, it requires a minimum amount of channel state feedback and still provides approximately a 30% improvement in throughput.

Extending Modulations, Coding Rates, and Frequency Bandwidth for High Throughput WLAN

Simulations show that by using adaptive bit and power-loading schemes, a significant performance increase over conventional modulation is achieved; but there is no increase beyond the maximum achievable data rate of 54 Mbps. Discussed below are two ways to enhance the maximum data rate in the context of single-antenna systems.

The first way is to use higher order modulations and/or higher coding rates to improve spectral efficiency. This technique puts more data into the channel. It would be analogous to speeding up traffic by having two cars share the same lane on the highway. Intel has investigated adaptive loading schemes in conjunction with an extended set of modulation types (including 256-QAM) and coding rates (R = 7/8). This approach can achieve a data rate of up

to 84 Mbps within a single 20 MHz channel which is the standard width for an 802.11a channel.

The second way is to use increased channel widths, which is analogous to widening the highway by adding more lanes. Simulations using up to four 20 MHz channels in one U-NII band in conjunction with the ABL schemes, showed approximately a 1.5x improvement over standard OFDM using the same number of channels. In an extended bandwidth mode, an ABL OFDM physical layer can achieve 300 Mbps, or approximately a 6x improvement, over single-channel standard OFDM.

Advanced FEC Coding

Nearly all wireless systems employ forward error correction (FEC) techniques to correct the numerous transmission errors that occur in wireless channels. The existing IEEE 802.11 wireless standards all specify a well-known convolutional FEC code and the Viterbi algorithm that is universally used for decoding these codes. At the time the current standards were developed, this approach was the most practical considering cost, complexity, power consumption, and decoding latency. Unfortunately convolutional codes and the Viterbi algorithm leave a significant amount of power efficiency unclaimed with respect to the theoretical capacity limits.

About ten years ago, the FEC community experienced a revolution with the discovery of Turbo codes, which were the first practical codes developed that could come close to theoretical capacity. Even after ten years of additional research, Turbo codes have seen limited acceptance in certain applications due to their complexity and decoding latency. One difficulty of Turbo codes is the performance degradation they experience when the size of the code block is reduced. For streaming applications such as broadcast video, this is not a problem, but WLAN packet sizes are often very small and transmissions as simple as packet acknowledgements require link reliabilities comparable to long packets. This degradation of performance for short blocks as well as their implementation complexity make it difficult to apply Turbo codes successfully to WLAN applications.

The discovery of Turbo codes, which use iterative decoding to achieve much of their performance, renewed research interest in another family of FEC codes known as low density parity check (LDPC) codes. Recent research has shown that carefully designed LDPC codes do not suffer the same performance degradation for short data blocks as Turbo codes. Intel has been researching a type of LDPC code for possible use in WLAN systems against the Viterbi decoder currently used in 802.11 systems.

Shortening the LDPC code for smaller blocks provides a significant improvement over the current system and gets us much closer to channel capacity. This approach results in increased range or higher throughput for systems utilizing this code.

High-Throughput MAC Considerations

Why do we need to modify the medium access control layer (MAC) at all? Isn't getting high throughput largely a physical layer (PHY) issue? The answer is "no" because while the PHY data rates are increasing significantly, the PHY overhead is not decreasing by the same factor. Therefore, throughput becomes dominated by overhead (radio turnaround times, modem pipeline delay, preamble, PLCP headers, and so forth.).

To make the most efficient use of the wireless medium, we need to take a system view and design the operation of the MAC so that it uses and manages the services of the PHY effectively while providing an unmodified interface to the higher layers. Some areas Intel is researching include investigations into how the MAC should manage burst packet traffic, quality of service (QOS) enhancements, channel feedback information and coexistence with legacy devices.

Conclusion

This paper has discussed techniques that can improve throughput and range for next-generation WLAN devices and the research Intel is doing in these areas. MIMO, adaptive OFDM techniques, extended modulations, wider bandwidths, advanced coding schemes and MAC improvements are likely to be used in next-generation WLAN products. These techniques can be dovetailed to provide a system that uses both MAC and physical layers to increases channel capacity and make the most of this capacity.

For more information on how Intel research is actively contributing to the development of higher-throughput WLANs and other advanced wireless technologies, consult www.intel.com/technology/wireless.